

Investigating the Performance of an Adiabatic Quantum Optimization Processor

Kamran Karimi, Neil G. Dickson, Firas Hamze, M.H.S. Amin, Marshall Drew-Brook,

Fabian A. Chudak, Paul I. Bunyk, William G. Macready, and Geordie Rose

*D-Wave Systems Inc., 100-4401 Still Creek Drive, Burnaby, B.C., Canada, V5C 6G9**

Adiabatic quantum optimization offers a new method for solving hard optimization problems. In this paper we calculate median adiabatic times (in seconds) determined by the minimum gap during the adiabatic quantum optimization for an NP-hard Ising spin glass instance class with up to 128 binary variables. Using parameters obtained from a realistic superconducting adiabatic quantum processor, we extract the minimum gap and matrix elements using high performance Quantum Monte Carlo simulations on a large-scale Internet-based computing platform. We compare the median adiabatic times with the median running times of two classical solvers and find that, for the considered problem sizes, the adiabatic times for the simulated processor architecture are about 4 and 6 orders of magnitude shorter than the two classical solvers' times. This shows that if the adiabatic time scale were to determine the computation time, adiabatic quantum optimization would be significantly superior to those classical solvers for median spin glass problems of at least up to 128 qubits. We also discuss important additional constraints that affect the performance of a realistic system.

I. INTRODUCTION

NP-hard optimization problems arise frequently in scientific and commercial applications [1, 2]. Their formal intractability, combined with the significant value that can be generated by finding better techniques for solving them, has led to a thriving research community searching for more effective and useful algorithms. While much progress has been made, for many applications the best current algorithms are not nearly good enough. An example of a very difficult optimization problem is the Eternity II puzzle [3].

Adiabatic Quantum Optimization (AQO) algorithms [4–8] may be effective tools to tackle hard optimization problems, making it important to study their performance potential. AQO algorithms are physics-inspired approaches that can be used as components of both exhaustive and heuristic

*Electronic address: kkarimi, ndickson, fhamze, mhsamin, marshall, fchudak, pbunyk, wgm, rose@dwavesys.com

solvers. In these algorithms, a Hamiltonian changes over time in such a way that after the computation is complete, the system has a non-zero probability of being in its lowest energy state, also known as its ground state. The ground state encodes the global optimum of the problem being solved. The excited states correspond to solutions that are not globally optimal. The M^{th} excited state corresponds to the $(M + 1)^{th}$ best solution.

AQO algorithms are not well understood. In particular it is very difficult to predict how their running time scales with problem size for any particular instance class of interest. Most of the work that has been done on this question focuses on the issue of whether, for asymptotically large problems, these algorithms require running times that scale polynomially or exponentially with problem size for some specific choice of solver parameters and instance class. Here we are not going to directly address these difficult issues, and approach the analysis of an AQO algorithm from a much simpler perspective.

Previous studies of the performance of adiabatic algorithms [5, 9, 10] were merely concerned with the scaling and not the actual value of the time in e.g., seconds for a realistic system. Moreover, in all those studies it is assumed, for example, that the temperature of a real system can always be kept much lower than the minimum energy gap between the ground state and first excited state, and that the adiabatic time is the only relevant timescale. Here, we present results suggesting that under the same assumptions, a superconducting hardware implementation [11] of an AQO algorithm could offer a speedup of several orders of magnitude over state of the art conventional approaches. However, we show that for a real processor with reasonable energy scales, the minimum energy gap is much smaller than feasible temperatures, even for fairly small problems, and thus, the simulations cannot accurately predict runtime scaling of such a processor. We also find that the adiabatic time is currently negligible compared to other necessary processes, such as readout of the results and thermalization (allowing the processor to cool down after a readout has been performed).

The rest of this paper is organized as follows. Section 2 presents relevant details of the simulated processor. Section 3 introduces the problems we solve and provides information about the techniques needed to perform our large-scale simulations. Section 4 presents the classical solvers used in this study and compares their performance with that of our simulated AQO processor. Section 5 concludes the paper.

II. THE SIMULATED PROCESSOR

In order to calculate adiabatic times for a quantum computation, in units of seconds, it is necessary to target a specific AQO processor architecture. In more theory-oriented publications, no limitation is placed on the connectivity of the architecture, but arbitrary connectivity becomes impractical to fabricate as the number of qubits increases. Here, we model a processor architecture with finite connectivity that has been fabricated. All of the simulation parameters come from measurements done on the processor in [11], so there were no free parameters. This processor is a programmable superconducting integrated circuit [12] designed to perform AQO with up to 128 pairwise-coupled [13] superconducting flux qubits [11]. To our knowledge, there are no other AQO processor architectures of comparable size under development.

For analyzing an AQO process of a closed system at zero temperature, the key details are captured in the hardware's Hamiltonian. As with the aforementioned studies, in order to perform the analysis, we make the assumption that the temperature of the processor is much smaller than the minimum energy gap, effectively zero, and that there is no disturbance from any outside environment. This approximation is clearly not exact, but is necessary, because of the infeasibility of simulating a non-trivial open system. The Hamiltonian we consider here is of the form

$$H(s) = A(s) H_I + B(s) H_P \quad (1)$$

H_I and H_P are dimensionless, and $s = t/t_f$ is a normalized time $0 \leq s \leq 1$. The initial dimensionless Hamiltonian as implemented by the hardware [11] is

$$H_I = - \sum_{i=1}^N \sigma_i^x \quad (2)$$

where σ_i^x is the Pauli X matrix for qubit i .

The final Hamiltonian, H_P , encodes the problem to be solved, and must be diagonal in the basis in which the qubits are read out. Because of this, the final state of the system will be classical, and therefore, qubits can be read out individually without destroying the state. The specific form we use is

$$H_P = \sum_{i=1}^N h_i \sigma_i^z + \sum_{i,j \in E} J_{ij} \sigma_i^z \sigma_j^z \quad (3)$$

where σ_i^z is the Pauli Z matrix for qubit i . The real-valued vector h and matrix J together define a *problem instance*, and are normalized such that $-1 \leq \{h_i, J_{ij}\} \leq +1$. E is the *allowed edge set*, dictating which J_{ij} can be non-zero (all other J_{ij} are zero). In other words, E is the set of pairwise

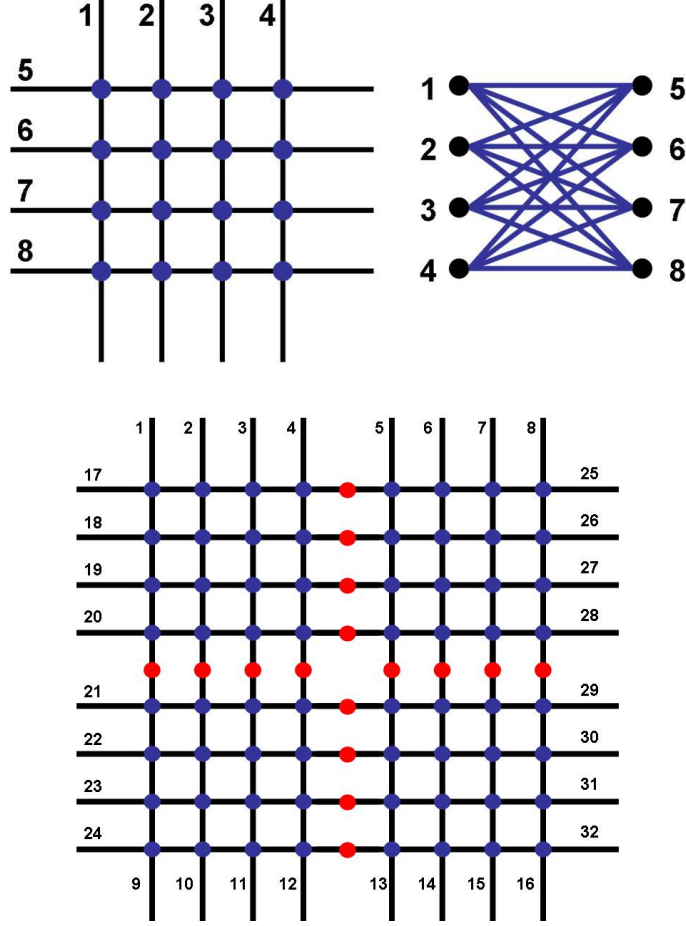


FIG. 1: (Colour online) Top: Two representations of the 8 vertex/unit cell graph (the complete bipartite graph on 4 vertices $K_{4,4}$). On the left, the variables are shown as black lines, with the allowed edge set E denoted by the set of blue points (intra-cell couplers) where the lines intersect. Note that this is not a conventional graph representation. On the right, showing the conventional graph representation, the variables are the black dots, and the non-zero edges are shown as the blue connecting lines. Bottom: 4 tiled $K_{4,4}$ unit cells, giving a total of 32 variables. Here the red circles represent edges between unit cells (inter-cell couplers). Larger edge sets are constructed by further tiling.

connections between qubits in our particular processor [11]. We restrict E to correspond to tiled bipartite $K_{4,4}$ graphs, as shown in Fig. 1.

The functions $A(s)$ and $B(s)$ in Eq. (1) have units of energy. Here we use the functional forms shown in Fig. 2. These functions are determined from the devices in the real processor we have simulated [11, 13–15], i.e. the processor uses a similar annealing schedule. The energy scales of $A(s)$ and $B(s)$ are limited by practical considerations for the actual superconducting devices used in the modeled processor, such as junction sizes, critical currents, etc. At the beginning of the

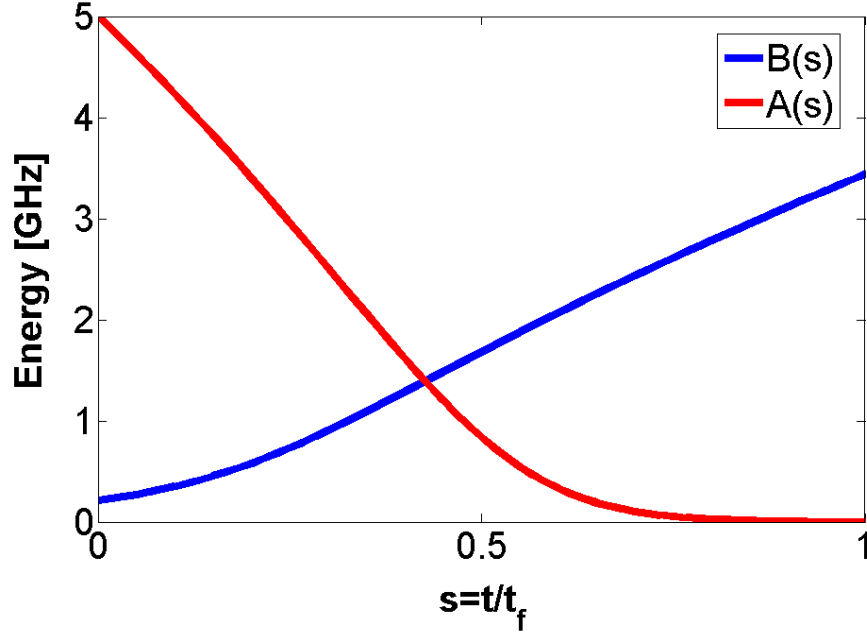


FIG. 2: (Colour online) Envelope functions $A(s)$ and $B(s)$ of the simulated processor.

computation (the leftmost side of Fig. 2 at $s = 0$), the transverse term (2) in the Hamiltonian (1) is dominant. As the computation proceeds, the strength of the diagonal term (3) grows while the transverse term shrinks. Once the transverse term has been turned off completely (the rightmost side of Fig. 2 at $s = 1$) the qubits are read out in the Z basis.

Note that while these are not the envelope functions $A(s) \sim (1 - s)$ and $B(s) \sim s$ used in other studies of AQO [5, 9], they do satisfy the underlying requirement of the AQO algorithms in that $A(0) \gg B(0)$ and $A(1) \ll B(1)$. We do not expect the results to change significantly if the linear envelope functions were used.

Changing s from 0 to 1 within time t_f at zero temperature results in finding the global optimum with measurable probability as long as $t_f \gtrsim t_a$, where [16, 17]

$$t_a = \frac{4 \left| \langle 1 | \frac{dH}{ds} | 0 \rangle \right|_{s=s^*}}{\pi g_m^2}, \quad (4)$$

which here we refer to as the adiabatic time scale. g_m is the minimum energy gap between the ground and first excited states of the Hamiltonian (1) and occurs at parameterized time s^* , which is unknown for any particular choice of $\{h, J\}$. Calculating t_a for a general, arbitrarily large problem may be very challenging, if at all possible. In such cases different values of t_f may be tried in order to obtain an acceptable rate of success.

TABLE I: The first three columns show the arrangement of unit cell graphs (shown in Fig. 1), and the resulting total number of variables in the examined problem instances. 8- and 32-variable instances are as shown in Fig. 1. Input to classical solvers were problems of these sizes. For simulations, the average number of s values for each problem size and the corresponding total number of variables in the simulated system are shown in the last two columns. The AQO simulations solved problems of these sizes. All simulations were performed with 256 Trotter slices.

# Rows	# Columns	# Variables	# s Values	Simulated Variables
1	1	8	30	61,440
2	1	16	40	163,840
2	2	32	50	409,600
3	2	48	60	737,280
3	3	72	70	1,290,240
4	3	96	110	2,703,360
4	4	128	130	4,259,840

III. CALCULATING THE ADIABATIC TIMES

We wish to calculate the performance of the AQO algorithm embodied in Eq. (1) on a group of problem instances ranging in size from 8 to 128 variables. Table I describes the tiling of bipartite $K_{4,4}$ graphs used to structure problem instances of different sizes.

It is necessary to choose a scheme for generating problem instances where these instances are all similar in some way. Here, we generate spin glass instances by randomly assigning all h_i and intra-cell couplers J_{ij} to be either $+1/3$ or $-1/3$ with equal probability. We set all inter-cell couplers to be $J_{ij} = -1$. We also require each instance to have a unique global minimum, so instances found to have multiple global minima were removed. Solving for the ground state of Eq. (1) at $s = 1$ using the $\{h, J\}$ generated using this prescription is an NP-hard optimization problem [18, 19]. For each problem size N , we randomly generated 100 problem instances. All problem instances used in this work are available online [20].

Calculation of the adiabatic time can be done by exact diagonalization for small problems (e.g. 8 or 16 variables), but not for larger problems, due to the exponential growth of the Hamiltonian. We use a discrete imaginary time quantum Monte Carlo (QMC) approach, as described in [9], to compute the minimum gap and matrix elements necessary for calculating the adiabatic time at all problem sizes considered in this paper. We apply the Suzuki-Trotter decomposition [21]

to be able to use classical Monte Carlo techniques (e.g. Metropolis) to perform the simulations. As described in [9], doing so introduces an effective temperature to the resulting system, which does not negatively affect the simulation, so long as this effective temperature is well below the minimum gap.

To perform our adiabatic time calculations, it was necessary to simulate the Hamiltonian (1) at different values of s (e.g. about 130 different values for the biggest problems we consider). Table I shows the average number of s values for problems of different size. For values of s close to 1, however, the resulting systems were such that single-site Monte Carlo algorithms suffered from severe equilibration issues. A straightforward *parallel tempering* (PT) [22] solution is to supplement the single-site dynamics with *exchange* moves of spin configurations between systems at adjacent values of s . This differs slightly from the way PT is usually applied to classical spin systems, where the simulations take place at different *temperatures*; in our context, the small and large s values play analogous roles to high and low temperatures respectively.

For systems with first-order phase transitions [10], it is not sufficient to merely ensure that adjacent systems' parameters are sufficiently close together to yield a reasonable (e.g. 25%) swapping probability. This problem was solved using an improved version of the *feedback-optimized PT* [23] technique, as explained in [24]. Without such a technique, obtaining reliable statistics for many of our systems would have been impossible.

In order to obtain the adiabatic times for AQO, each problem instance was simulated about 100 times, for a total of about 70,000 (7 problem sizes \times 100 instances per size \times 100 solutions per instance) simulations. For a typical 128-qubit problem, we needed to perform millions of Monte Carlo sweeps on about 4,259,840 (128 qubits \times 256 Trotter slices \times 130 s values) variables, which required huge computing resources. We used multi-threading to effectively use multi-core CPUs and speed up the execution of the simulation program [25]. To increase the performance of each thread, we vectorized the core of the Monte-Carlo method using Streaming SIMD Extensions 2 (SSE2) instructions. As a result of these and other optimizations, on an Intel Core i7-965 running at 3 GHz, the performance of our application increased about 50-fold compared to the single-threaded, unoptimized version.

We used the Berkeley Open Infrastructure for Network Computing (BOINC) [26] as a distributed computing platform. The resulting BOINC project, AQUA@home (Adiabatic QUantum Algorithms) [27], deployed our simulation application onto approximately 3,500 volunteer computers with a total of about 8,000 cores for around 3 months. Approximately 10^{20} floating point and integer operations were performed during the course of these computations. The simulation

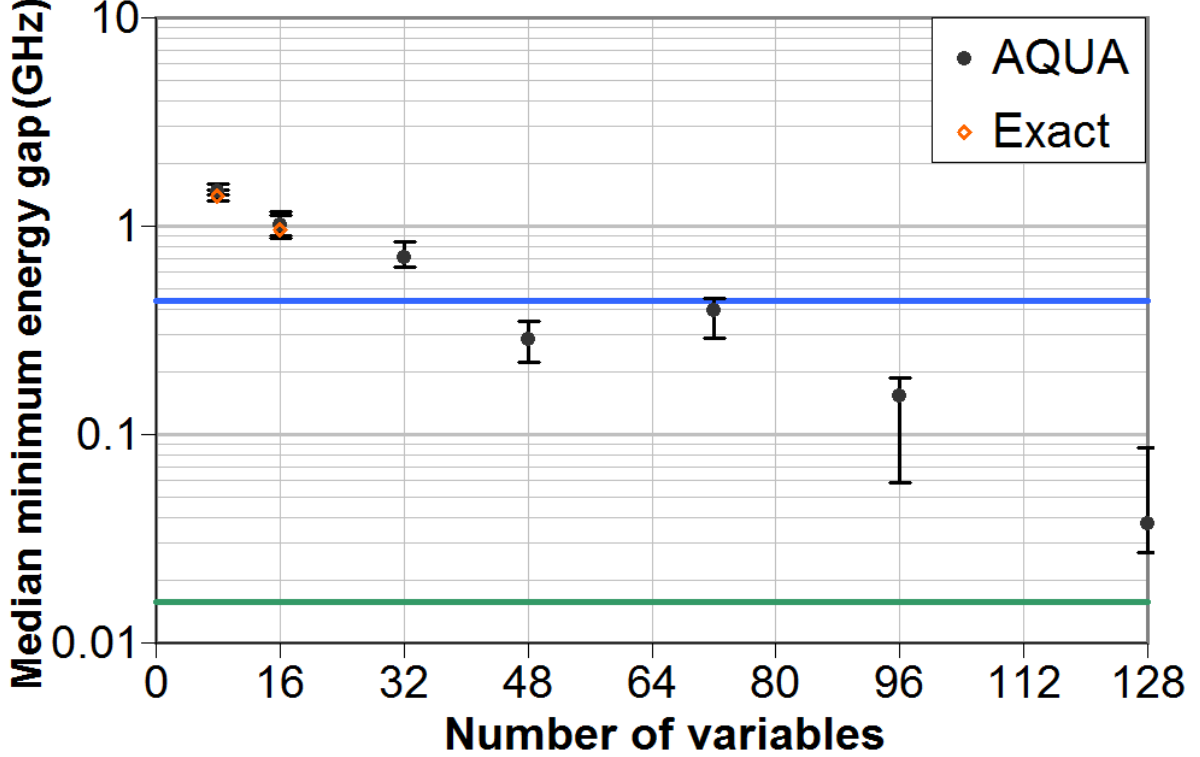


FIG. 3: (Colour online) Median minimum energy gaps as a function of problem size. Each data point is the median of the minimum energy gaps of 100 randomly generated instances. Error bars represent the 40% and 60% intervals around the median. The green line shows the simulation temperature of 0.75 mK (0.015625 GHz). The blue line shows the lowest feasible temperature of the real processor, 21 mK (0.44 GHz).

results, needed for computing Eq. (4), are available online for all 700 instances studied [28]. We verified our simulations for 8- and 16-variable problems by comparing the simulation results with exact values obtained from diagonalization.

Similar to other AQO publications [5, 9, 10], we use median values in our performance evaluations. The main reason is that the median of minimum energy gaps and running times are more useful than worst-case values, as they apply to typical problems. Notice that problems specifically designed to be hard (such as encryption) fall beyond our analysis. Another important reason for using median values is that the Quantum Monte Carlo (QMC) simulation technique we used is not able to determine the numerical values of extremely small energy gaps, making it nearly impossible to investigate worst-case problems.

In Fig. 3 we plot the median minimum energy gaps as a function of problem size. The simulation temperature of 0.75 mK (0.015625 GHz) is shown in green. We see that even for the 128-variable problems, the median minimum energy gaps are larger than this temperature. This implies that

for the considered problems, the extraction of ground state properties from the QMC simulations is reliable. The appropriateness of using the median values is confirmed by considering the error bars on the gathered data. We observe that running 100 instances provides a tight clustering of minimum gap values and running times.

Fig. 3 appears to show that the median gap size for the 72-qubit instances is larger than that of the 48-qubit instances. However, due to the nonzero error bars, this observation may not be correct. Also, while it is very common for median gap sizes to decrease overall with respect to the problem size, there is no evidence to suggest that for these or similar problem instances, the median gap size should always decrease monotonically with system size.

In the same figure, we plot (blue line) the lowest physical temperature, i.e., 21 mK, that can be realistically achieved by the processor under normal conditions in a dilution refrigerator. As can be seen, the median gap starts to get below the temperature at around 48 variables. Even for the smallest sizes, the median gap is only a factor of 2 to 3 larger than T . Therefore, the assumption of $T \ll g_m$, commonly made in most calculations related to adiabatic quantum computation is not realistic.

IV. COMPARISON WITH CLASSICAL SOLVERS' PERFORMANCE

We compare our AQO implementation's adiabatic running time performance with that of two conventional solvers: IBM's CPLEX [29], and MadCat. CPLEX is a well-known commercial solver for discrete optimization problems, and using it provides a baseline for a general optimizer's performance. MadCat is a complete solver developed in-house, and designed to exploit the known structure of problem instances. MadCat is a particular case of a well-known general algorithm [30] and relies on a *tree decomposition*: a rooted tree whose nodes are subsets of the problem's variables, subject to certain conditions. Optimal assignments of variables in a given node can be determined, conditioned on the values of variables in the parent node. At the root node, unconditional optimal assignments are found and this information is propagated out to the leaves to produce a full optimal solution. The size of the largest tree node, minus 1, is called the *treewidth* of the decomposition. The dynamic programming approach reduces MadCat's running time to be exponential in the treewidth, rather than exponential in the total number of problem variables. The $M \times N$ tiled graphs considered here have treewidth $\min(4M, 4N)$, so the small instance sizes we consider here can be solved quickly. We also used MadCat to test the non-degeneracy of the problem instances.

CPLEX searches the space of solutions less efficiently than MadCat. To speed-up CPLEX's

searches, we solved each instance using MadCat and obtained the optimal solution’s energy. We then ran CPLEX with that energy as an argument to indicate that it should stop as soon as it finds a solution with that energy, i.e. as soon as it finds the optimal solution.

Fig. 4 shows the computed median adiabatic times, as well as the median running times required by both CPLEX and MadCat to find the global optimum of the same 700 instances. Both classical solvers were run on a system with 2 quad-core Xeon E5430 2.66 GHz processors running 64-bit Linux. CPLEX was set to use all 8 cores. Each instance was solved 10 times for both solvers, and of those, the minimum running time was retained to make sure that non-deterministic behaviour (for CPLEX), or transient system loads (for CPLEX and MadCat) did not negatively impact the measured running times.

Using different classical solvers, or different computer hardware for running them, would provide different results. However, since we already use reasonably high-performance software and hardware, we do not expect such changes to substantially affect the paper’s conclusions.

It is important to emphasize that the predictions of running times for the AQO processor are not valid for problems with 48 or more variables running on the real processor at 21 mK, the lowest feasible temperature for the system. This is because the median minimum gap size drops below this operating temperature, as shown in Fig. 3. It is unknown, and non-trivial to determine, whether the true performance would be better or worse than the performance of a hypothetical processor at 0.75 mK [17]. If it is feasible to construct hardware with energy scales large enough that these median minimum gaps are increased to be larger than the operating temperature, such a processor would have significantly shorter adiabatic times than in Fig. 4, because of Eq. (4).

V. CONCLUDING REMARKS

This paper provides adiabatic running times for a specific AQO processor using distributed high-performance simulations, and compares the results with those of two classical solvers, i.e., CPLEX and MadCat. For the largest instances studied, MadCat is roughly 2 orders of magnitude faster than CPLEX. This improvement is because MadCat was designed to take advantage of the structure in the allowed edge set. Note that the median running times for both solvers grow exponentially for larger problems of this type. More striking is the computed median adiabatic time for the superconducting adiabatic quantum processor, which is approximately 4 orders of magnitude shorter than that of MadCat. So, despite not knowing the performance scaling of AQO for larger problems, we have shown that under the assumption that adiabatic time correctly represents the

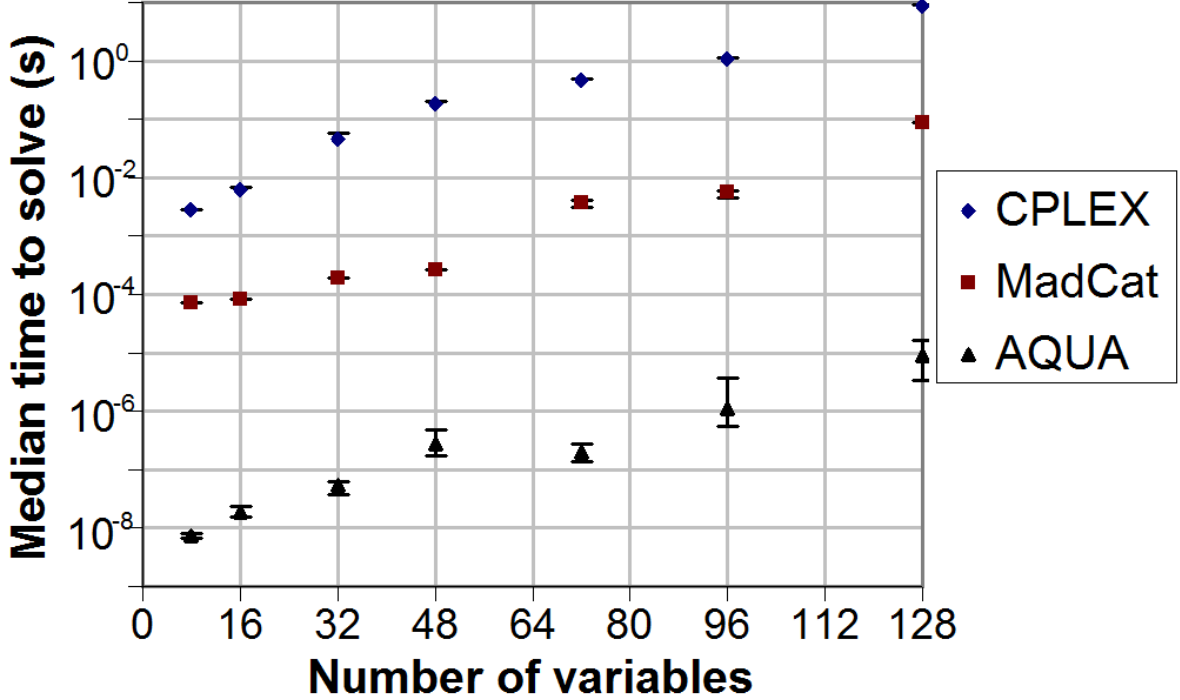


FIG. 4: (Colour online) Black: Median adiabatic times as a function of problem size. Each data point is the median of the adiabatic times of 100 instances. Red (Blue): Median times to find the global optimum for MadCat (CPLEX). Error bars represent 40% and 60% intervals around the median.

time scale for computation, an AQO processor could significantly outperform classical solvers.

Two important points have to be noted here. First, the above consideration ignores the effect of the environment, especially the fact that the lowest feasible temperature at which such a processor can be operated is larger than the median minimum energy gaps of problems with 48 or more variables. Although recent calculations suggest that a weak coupling to the environment does not significantly affect the time required to reach the final ground state with a measurable probability, even if the minimum gap is below the temperature [17, 31–37], a fair comparison should consider the performance of an open and not closed system. Unfortunately, such open system simulations are not feasible beyond ~ 20 qubits [17].

The second point is that for the small problems we investigate here, the adiabatic time does not dominate the running time of the real hardware. A fair comparison between the classical solvers and the AQO processor should include the time needed for operations such as programming the chip, readout, and thermalization (to ensure that the chip returns to its operational temperature before the next problem is solved). In the real processor, readout and thermalization times completely dominate the problem solving time for these small-scale problems. At the time of writing this

paper, serial readout takes roughly $36\ \mu s$ per qubit, and thermalization time is experimentally chosen to be $1000\ \mu s$, as compared with median adiabatic time of $10\ \mu s$ for 128 variable problems. These numbers depend on the processors' design and fabrication details, as well as the efficiency of cryogenic components used to cool the processor to 21 mK, and are expected to change over time.

Acknowledgments

The authors thank Edward Farhi, Helmut Katzgraber, and A. Peter Young for fruitful discussions. We would also like to enthusiastically thank the volunteer community supporting the AQUA@home distributed computing project, and the BOINC development team at UC Berkeley for their support.

References

-
- [1] Garey M K and Johnson D S, 1997, *Computers and Intractability, A Guide to the Theory of NP-Completeness*, W.H. Freeman & Co., New York
 - [2] Woeringer G J, 2003, Exact Algorithms for NP-Hard Problems: A Survey, *Lect. Notes Comput. Sci.* 2570 pp.185-207
 - [3] <http://us.eternityii.com/>
 - [4] Kadowaki T and Nishimori H, 1998, Quantum annealing in the transverse Ising model, *Phys. Rev. E* 58, 5355
 - [5] Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A, and Preda D, 2001, A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem, *Science* Vol. 292, no. 5516, pp. 472 - 475
 - [6] Das A and Chakrabarti B K, 2008, Colloquium: Quantum annealing and analog quantum computation, *Rev. Mod. Phys.* 80, 10611081
 - [7] Amin M H S, 2008, Effect of Local Minima on Adiabatic Quantum Optimization, *Phys. Rev. Lett.* 100, 130503
 - [8] Amin M H S and Averin D V, 2008, Macroscopic Resonant Tunneling in the Presence of Low Frequency Noise, *Phys. Rev. Lett.* 100, 197001
 - [9] Young A P, Knysh S, and Smelyanskiy V N, 2008, Size Dependence of the Minimum Excitation Gap in the Quantum Adiabatic Algorithm, *Phys. Rev. Lett.* 101, 170503

- [10] Young A P, Knysh S, and Smelyanskiy V N, 2010, First order transition in the Quantum Adiabatic Algorithm, *Phys. Rev. Lett.* 104, 020502
- [11] Harris R, *et al*, 2010, Experimental Demonstration of a Robust and Scalable Flux Qubit, *Phys. Rev. B* 81, 134510
- [12] M. W. Johnson *et al.*, 2010, A scalable control system for a superconducting adiabatic quantum optimization processor, *Supercond. Sci. Technol.* 23, 065004
- [13] Harris R *et al*, 2009, Compound Josephson-junction coupler for flux qubits with minimal crosstalk, *Phys. Rev. B* 80, 052506
- [14] Lanting T *et al*, 2009, Geometrical dependence of the low-frequency noise in superconducting flux qubits, *Phys. Rev. B* 79, 060509
- [15] Harris R *et al*, 2009, Synchronization of multiple coupled rf-SQUID flux qubits, *New J. Phys.* 11, 123022
- [16] Childs A M, Farhi E, Goldstone J, and Gutmann S, 2002, Finding cliques by quantum adiabatic evolution, *Quantum Inf. Comput.* Vol. 2, No 3 pp. 181-191
- [17] Amin M H S, Truncik C J S, and Averin D V, 2009, Role of single-qubit decoherence time in adiabatic quantum computation, *Phys. Rev. A* 80, 022303
- [18] Baharona F, 1982, On the computational complexity of Ising spin glass models, *J. Phys. A: Math. Gen.* 15, pp. 3241-3253
- [19] Istrail S, Statistical Mechanics, Three-Dimensionality and NP-completeness: I. Universality of Intractability for the Partition Function of the Ising Model Across Non-Planar Lattices, 2000, *32nd ACM Symposium on the Theory of Computing (STOC'00)*, ACM Press, pp. 87-96
- [20] <http://aqua.dwavesys.com/problems.html>
- [21] Suzuki M, 1976, Relationship between d-Dimensional Quantal Spin Systems and (d+1)-Dimensional Ising Systems, *Progress of theoretical physics* Vol. 56, No. 5, pp. 1454-1469
- [22] Hukushima K and Nemoto K, 1996, Exchange Monte Carlo Method and Application to Spin Glass Simulations, *J. Phys. Soc. Jpn.* 65 pp. 1604-1608
- [23] Katzgraber H G, Trebst S, Huse D A, and Troyer M, 2006, Feedback-optimized parallel tempering Monte Carlo, *J. Stat. Mech: Theory Exp.* 2006(03):P03018
- [24] Hamze F, Dickson N G, and Karimi K, 2010, Robust Parameter Selection for Parallel Tempering, *Int. J. Mod. Phys. C* Volume: 21, Issue: 5 pp. 603-615
- [25] Karimi K, Dickson N G, and Hamze F, 2010, High-Performance Physics Simulations Using Multi-Core CPUs and GPGPUs in a Volunteer Computing Context, *Int. J. High Perform. Comput. Appl.*, doi:10.1177/1094342010372928
- [26] Anderson D P, 2004, BOINC: a system for public-resource computing and storage, *Fifth IEEE/ACM International Workshop on Grid Computing (GRID'04)*, pp. 4-10
- [27] <http://aqua.dwavesys.com/>
- [28] <http://aqua.dwavesys.com/results/>
- [29] <http://www.ilog.com/products/cplex/>

- [30] Kask K, Dechter R, Larrosa J, and Fabio G, 2001, Bucket-tree elimination for automated reasoning, *Artif. Intel.* 125, pp. 91, 131.
- [31] Childs A M, Farhi E, and Preskill J, 2001, Robustness of adiabatic quantum computation, *Phys. Rev. A* 65, 012322.
- [32] Roland J and Cerf N J, 2005, Noise resistance of adiabatic quantum computation using random matrix theory, *Phys. Rev. A* 71, 032330.
- [33] Tiersch M and Schützhold R, 2007, Non-Markovian decoherence in the adiabatic quantum search algorithm, *Phys. Rev. A* 75, 062313.
- [34] Ashhab S, Johansson J R, and Nori F, 2006, Decoherence in a scalable adiabatic quantum computer, *Phys. Rev. A* 74, 052330.
- [35] Amin M H S, Love P J, and Truncik C J S, 2008, Thermally Assisted Adiabatic Quantum Computation, *Phys. Rev. Lett.* 100, 060503.
- [36] Wan A T S, Amin M H S, and Wang S X, 2009, Landau-Zener Transitions in the presence of spin environment, *Int. J. Quantum Inf.* Vol. 7, No. 4, pp. 725-737.
- [37] Amin M H S, Averin D V, and Nesteroff J A, 2009, Decoherence in adiabatic quantum computation, *Phys. Rev. A* 79, 022107.